

(19)



Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) Publication number:

0 239 298
A2

(12)

EUROPEAN PATENT APPLICATION

(21) Application number: 87302190.1

(51) Int. Cl.: G 06 F 13/22

(22) Date of filing: 13.03.87


(30) Priority: 26.03.86 US 844911

(71) Applicant: **TANDEM COMPUTERS INCORPORATED,**
19333 Vallco Parkway, Cupertino California 95014 (US)(43) Date of publication of application: 30.09.87
Bulletin 87/40(72) Inventor: **Fuller, William T.,** 1536 Shasta Avenue, San
Jose California 95126 (US)

(84) Designated Contracting States: DE FR GB IT SE

(74) Representative: **Huntingford, David Ian et al, W.P.**
THOMPSON & CO. Coopers Building Church Street,
Liverpool L1 3AB (GB)(54) **Modifiable input/output priority selection.**

(57) An apparatus which determines the order of data communication between a plurality of peripheral devices that wish to do so and a central processor unit. Determination is made according to one of a number of selectable priority schedules. The apparatus is modifiable by programmed control so that certain of the peripheral devices can have their priorities re-configured depending upon changing circumstances.



EP 0 239 298 A2

DESCRIPTIONMODIFIABLE INPUT/OUTPUT PRIORITY SELECTION

5 The invention is directed generally to data communications between a processor unit and a number of peripheral devices of a data processing system. More particularly, the invention relates to a modifiable input/output priority selection apparatus that permits
10 the processor unit to serve a plurality of peripheral devices desiring to communicate with the processor unit by determining the order of communication.

 Many of today's data processing systems include a central processing unit that is coupled to a variety
15 of peripheral devices that store data, receive data from remote sources, or (in the case of terminal and keyboard units, for example) provide a communication interface between a user and the processing system itself. Often, the peripheral devices communicate over a
20 single bus system of one sort or another, so that in order to avoid data communication collisions, when two or more of the peripheral devices require interaction with the processor unit, the ordering of such interaction or communication with the separate devices must
25 be established.

 This ordering, however, must also take into account the speed with which different types of peripheral devices can communicate data, and the permanency of the data to be communicated (i.e., how long will the
30 data be available before lost if not transmitted). For example, control devices for communications equipment sometimes will have data buffers that are very small (two bytes). When a communications control device wants to transfer data to the processor unit, it must do so
35 almost immediately. Accordingly, the ordering determination is often done according to various priorities

assigned to the peripheral devices. These priority assignments specify which peripheral device receives attention first if more than one requires the processor unit at a given time.

5 In one form of prioritization heretofore used, peripheral device priorities are based upon the position of the peripheral device along a daisy-chained input/output (I/O) bus. Changing the assigned priority of a peripheral device, therefore, requires physical relocation
10 of the peripheral device along the I/O bus.

 Yet another prioritization technique relies upon switch settings in the peripheral device which set the priority of that device, relative to other devices coupled to the same I/O bus. Again, although no
15 physical relocation of the peripheral device is necessary to modify its priority in the hierarchy of prioritization, a physical act must still be taken to reset the switches if the assigned priority of a peripheral device is to be changed.

20 Thus, while various prioritizing techniques are known, none provides any real flexibility in changing assigned priorities without some form of system reconfiguration by physical means.

 It is an object of the present invention to
25 improve on the foregoing known techniques.

 The present invention, provides apparatus that is modifiable by programmed control to selectively reconfigure the prioritization of data communications between a number of peripheral devices and a central
30 processing unit. In addition, the apparatus provides automatic readjustment of the prioritization to give certain peripheral devices higher priority than others for one period of time, and then restructuring to grant other peripheral devices higher priorities than
35 before.

 In its broadest aspect, the present invention comprises a presettable register for containing

information used to select one of a number of available predetermined priority schedules; each schedule assigns data communication priorities to a number of peripheral units forming a part of a data processing system. Logic
5 is provided for combining the information with data that indicates which of the peripheral devices desire communication with the processor unit, producing therefrom a selection signal that identifies the highest priority peripheral device desiring data
10 communication. After selection of the peripheral is made, the presettable register can be modified to select another of the predetermined priority schedules, or left to specify the same priority schedule, as desired.

15 The presettable register is accessible to the processor unit. This allows program control over priority schedule selection.

The invention is described further hereinafter, by way of example only, with reference to the
20 accompanying drawings, in which:

Fig. 1 is a representative illustration of a data processing system that includes at least one processor unit coupled to a plurality of device controllers that, in turn, each control one or more peripheral
25 devices (not shown);

Fig. 2 is a block diagram that illustrates generally the priority selection apparatus of the present invention; and

Fig. 3 is a block diagram of one of the sub-group
30 priority selection devices of the invention of Fig. 2.

The present invention is used in a multiprocessor system of the type disclosed in U.S. Patent No. 4,228,496, and is an improvement of the prioritization technique used therein. Generally, some form of priori-
35 tization schedule is needed when two or more peripheral devices desire communication with a processor unit.

Certain peripheral devices (such as, for example, a communications controller, which may have a very small data buffer and, therefore, the need to empty that buffer promptly) require a higher frequency of attention than other devices (such as, for example, keyboards and terminals). Accordingly, higher priority is usually given to those devices that must unload data quickly to the high-speed devices, and the like, while a lower priority is given to slower devices.

Illustrated in Fig. 1 is a data processing system, designated generally with the reference numeral 10, comprising at least one processor unit 12 coupled to a plurality of device controllers 14 by an input/output (I/O) bus 16. Although not specifically shown, it will be understood by those skilled in this art that each device controller 14 will have attached to it peripheral devices (not shown) such as disk drives, tape units, terminals, data communications equipment (i.e., modems), and the like. Each device controller 14 is responsible for selecting one of its peripheral devices for communication with the processor unit 12 when required.

Generally, data communications are initiated in the first instance by the processor unit 12 when it desires to retrieve information from a peripheral device (not shown), or wishes to store information thereat. Accordingly, a request signal is sent by the processor unit 12, via the I/O bus 16, to the device controller 14 associated with the peripheral device (not shown) that has the desired information, or is to be the location at which information will be stored. The device controller 14, in turn, goes through a process of setting up the peripheral device (not shown) for the communication, signalling the processor unit 12 when ready to transfer data.

It is not uncommon to have a number of device controllers 14 ready to transfer at any one time in

response to data transfer requests that have been made by the processor unit 12. The processor unit 12 is then faced with the task of determining which of the device controllers 14 desire communication, and arranging the communication colloquy in an orderly fashion, preferably pursuant to some form of prioritization schedule. For a further discussion of how data communications between processor units and device controllers 14 via the I/O bus 16 may be handled, the reader is referred to U.S. Patent No. 4,228,496, the disclosure of which is incorporated herein by reference. Determination of the order of communications is the concern of the present invention, and Fig. 2 shows preferred apparatus for determining such order.

Referring now to Fig. 2, a modifiable input/output priority selection apparatus, designated generally with the reference numeral 20, is shown as receiving a portion 16a of the I/O bus 16 (Fig. 1) at a data-in register 22. The I/O bus portion 16a includes a 16-bit, bidirectional data bus that communicates data between the processor unit 12 and the device controllers 14, according to one of a number of prioritization schedules maintained by the priority selection apparatus 20. Also coupled to the I/O bus portion 16a is a 16-bit data-out register 24. The I/O bus portion 16a is also coupled to the circuitry (not shown) in the processor unit 12 for actually handling the data transfers between the processor unit and the device controllers 14 (using the data-in and data-out registers 22, 24).

Data-in and data-out registers 22, 24 are used for polling and selection. In the manner described in more detail in the above-referenced U.S. Patent No. 4,228,496, each bit position of the data-in register 22 corresponds to a specific device controller 14. During polling, those device controllers 14 requiring or requesting a data transfer between it and the processor

unit 12 will set its corresponding bit position in the data-in register 22 to a ONE. At the completion of a poll one or more of the 16 available bit positions may
5 be set to a ONE, indicating that the device controllers 14 are ready for data communication.

The content of the data-in register 22 is coupled, in four-bit groups, to subgroup priority selection circuits 26, 28, 30 and 32. If two or more of the
10 bit positions of each four-bit group are set to a ONE, the corresponding subgroup priority selection circuits 26-32 determine from this information, according to a prearranged priority, which device controllers corresponding to the set bit positions represent shall be
15 selected to communicate first.

An encoded representation of the selected device controller 14 is coupled from the subgroup priority selection circuits 26-32 to a multiplex circuit (MPX) 34. In addition, each subgroup priority selection
20 circuit 26-32 couples an available (AVAIL) signal, indicating that one or more of the device controllers 14 associated with the particular subgroup priority selection circuit is ready, to a master priority selection circuit 40. The master priority selection circuit 40
25 functions, in turn, to select among the subgroup priority selection circuits 26-32 as to which subgroup of device controllers 14 shall proceed with data communication, in much the same manner that the subgroup priority selection circuits operate. The master priority selection
30 circuit 40 produces a SELECT signal that is applied to the selection inputs of the MPX 34 to select which of the encoded outputs of the subgroup priority selection circuits 26-32 are to be coupled to a 4-to-16 decode circuit 42.

35 The SELECT signal from the master priority selection circuit 40 and the output of the MPX 34 form a complete identification of which of the device controllers 14 have been selected, in encoded form. This

4-bit encoded designation is applied to, and decoded by, the 4-to-16 decoder 42, when enabled by an ENABLE signal provided by the master priority selection circuit 40. The 16-bit output of the 4-16 decoder 42 is coupled to the data-out register 24, and used to set one bit position of the data-out register 24, corresponding to the selected device controller. The processor unit 12 will then communicate the contents (15 ZEROS and one ONE) data-out register 24 to the data bus 16a. The device controller assigned to the bit-position of the data-out register 24 is thereby informed that it is selected.

The SELECT and ENABLE signals generated by the master priority select circuit 40 are communicated to a 2-to-4 decode circuit 44. The ENABLE signal initiates (enables) operation of the 2-4 decode circuit, while the SELECT signal, which identifies the selected device controller subgroup, is used to produce an enable signal (EN.0-EN.3) that is used to modify the priority schedule of the corresponding subgroup priority selection circuit, as will be described below. The signals EN.0-EN.3 are respectively coupled to the subgroup priority selection circuits 26-32. Only one of the signals EN.0-EN.3 will be active at any one time.

Before discussion of the architecture of the individual subgroup priority selection circuits 26-32 and the master priority selection circuit 40, it may be beneficial to understand the role played by the priority selection apparatus 20 in prioritizing communications between the processor unit and two or more of the device controllers 14. A data communication on the I/O bus 16 is initiated when a data transfer command is sent by the processor unit 12 to a one of the device controllers 14 (to which is connected the peripheral unit (not shown) with which the communication is involved). The processor unit 12 then resumes normal processing activity, awaiting a "Ready" signal from the

designated device controller 14, indicating that it is ready to send or receive data, depending upon the request made by the processor unit 12.

5. At times, a number of requests may be issued by the processor unit 12 so that often a number of Ready indications may be pending and requiring response by the processor unit 12. In these cases the processor unit 12 must first determine which of the device controllers 14 are requesting service, and if there are two or more, what their priorities are. Accordingly, the processor unit 12 will issue commands on the command lines of I/O bus 16 to "poll" the device controllers 14, causing each device controller 14 to apply a ONE on the data line of the I/O bus portion 16a to which that particular device controller 14 is assigned. The content of the I/O bus portion 16a is set in the data register 22, and the "polling" sequence is terminated. Thus, at the end of the poll, initiated by the processor unit 12, the content of the data-in register will contain an indication of which of the device controllers 14 are ready for communication.

The subgroup priority selection circuit 26 is illustrated in greater detail in Fig. 3. The architecture, function and operation of the subgroup priority selection circuits 28, 30 and 32 are substantially identical in nature, so that any discussion of the subgroup priority selection circuit 26 should be understood as being applicable to the subgroup priority selection circuits 28-32 equally. Substantially the same can be said about the master priority selection circuit 40; although there are some differences, the design and function of the master priority selection circuit is basically the same as any of the subgroup priority selection circuits 26-32. Accordingly, the following description of subgroup priority selection circuit 26 should be taken to apply equally to the subgroup and

master priority selection circuits 28-32 and 40, except as noted.

Referring now to Fig. 3, the subgroup priority selection circuit 26 is illustrated as receiving
5 its 4-bit subgroup information on requestor lines R0, R1, R2 and R3 from the data-in register 22 (Fig. 2) at a selection logic unit 50. The 4-bit subgroup signal lines are also applied to an OR gate 52 producing the
10 AVAIL signal that is communicated to the master priority selection circuit 40.

Also applied to the selection logic unit are the contents of a 3-bit state register 54 and a 2-bit mode register 56. These two registers are presettable
15 from data that is communicated from a five-bit portion of an internal data bus 58 from the CPU (not shown) of the processor unit 12. Three bits of the internal data bus 58 are coupled to the state register 54 through a multiplex circuit (MPX) 60, while the remaining two
20 bits of the internal data bus 58 are coupled directly to the mode register 56. The state register 54 also receives three bits of NEXT STATE information from the selection logic unit 50. Control of the MPX 60 to select the source of data to be set in the state register
25 54 is effected by a LOAD signal produced by the CPU (not shown). When present, the LOAD signal routes the three-bit portion of the internal data bus 58 to the data input of the state register. When not present the MPX 60 defaults to selection of the NEXT STATE information.
30

The LOAD signal is padded through an OR gate 61 to an AND gate 62, and used to gate a CLK signal to set the output of the MPX 60 into the state register 54. The state register 54 is also enabled, for load-
35 ing, by the EN.0 signal from the 2-to-4 decode circuit 44 (Fig. 2).

The state register 54 defines which one of eight available prioritization schedules are to be used.

In addition, the capability of periodically presetting the state register 54 (either with the NEXT STATE information or data on the bus 58) allows the prioritization schedules to be cycled through in various orders. The mode register 56 defines the particular order of cycling, or whether or not any cycling will be used at all. Table I; below, sets forth the priority schedules available, showing the "Current State" content of the state register 54 used to select one of the available eight priority schedules, and the "Next State" that the state register 54 will be set to (by the NEXT STATE information) when a state transition is made.

TABLE I

MODE = 11			MODE = 01		
CURRENT STATE	NEXT STATE	PRIORITY SCHEDULE	CURRENT STATE	NEXT STATE	
0	1	R0 R1 R2 R3	0	1	
1	2	R1 R2 R3 R0	1	3	
2	3	R0 R1 R2 R3	-	-	
3	4	R2 R3 R0 R1	3	7	
4	5	R0 R1 R2 R3	-	-	
5	6	R1 R0 R2 R3	-	-	
6	7	R0 R1 R2 R3	-	-	
7	0	R3 R0 R1 R2	7	0	

Table I illustrates two modes of operation, each determined by a state of mode register 56. The two left-hand columns (Current State - Next State) define a mode of operation in which the state register cycles through its eight states in sequential fashion. This mode is selected by a "11" in mode register 56. The two right-hand columns show the state transitions (Current State - Next State) of state register 56 when the mode register 56 is a "01." The center column (labeled Priority Schedule) shows the available priority schedules for the associated bit positions

(communicated on the requestor lines R0-R3) of the data-in register 22 for the corresponding current state.

Table I is interpreted as follows: Reading
 5 each row from left to right, the first row (Current State) shows the content of state register 56, the state identification (Next State) the state register 56 will be set to by the EN.0 signal (and CLK), and the priority order (Priority Schedule) of the bit positions on the
 10 incoming requestor lines R0-R3 (highest priority on the left). Continuing, the next two rows (Current State, Next State) described, as noted above, the state transitions of the state register 56 when the mode register 56 is a "01.". An example may help in understanding
 15 Table I.

Assume the following conditions: The Current State of state register 54 is "00"; the mode register 56 is "11"; the bit locations for incoming requestor lines R1 and R2 of the data-in register 22 are a ONE while
 20 the remainder of the data-in register 22 is ZERO. Given these conditions, the selection logic unit 50 will produce a two-bit coded signal that is applied to the MPX 34, identifying the device controller 14 associated with the bit position brought in by requestor line R1
 25 as the one determined by the subgroup priority selection circuit 26 to communicate with the processor unit 12; the device controller 14 associated with the bit position of requestor line R2 must wait its turn.

When this communication is complete, the EN.0
 30 signal is generated. At the same time the selection logic unit 50 derives NEXT STATE information from the content of state and mode registers 54 and 56, producing a "1" that is set in the state register 56. This new content of the state register 56 now defines the
 35 Current State: "1." The new Priority Schedule, R1 (highest priority), through R0 (lowest priority), is defined.

Another poll will be conducted, causing bit position for requestor R1 to be set to a ZERO, since it has just been serviced, and setting bit position brought in-on requestor line R2 again to a ONE (since the associated device controller's request for communication is yet to be honored). The output of the priority selection unit 50 proposed an R2 selection and, assuming no other bit position selections of higher priority, a ONE is set in the data-out register 24 to signal the device controller 14 associated with bit position R2 that can now communicate with processor unit 12. The next state register 54 is set to identify a new Current State:

"2."

15 If the mode register 56 is set so that $MODE = 01$, the priority schedule selections are in a "round-robin" manner; that is, the procession of states of the state register 54 is "0"- "1"- "3"- "7"- "0" and so on. As with $MODE = 11$, transitions from the current state to the next state in $MODE = 01$ occur after the device controller 14 associated with the current selection bit position on incoming requestor lines R0-R3 is selected.

25 If the least significant bit position of the mode register is a ZERO ("0"), the current state will be frozen and no "next state" transition will take place. Thus, if the mode register 56 is a "10" or a "00," the current state becomes the next state perpetually, until the least significant bit of the mode register becomes 30 a "1" in which the most significant bit dictates the type of current state to next state transition as described above.

As indicated above, the subgroup priority selection circuits 28-32 are identical in structure and operation to the subgroup priority selection circuit 35 26. Only the bit positions of the data-in register 22 received by each are different.

Similarly, the structure and operation of the master priority selection circuit 40 is the same, with the following exceptions. In place of the bit positions brought in by the requestor lines R0-R3, the master priority selection circuit 40 receives the AVAIL signals from each of the subgroup priority selection circuits 26-32, selecting one of them, according to a priority schedule, for use as the SELECT signal. Also, the state register (not shown) of the master priority selection circuit 40 is changed in the same manner as that described with respect to the subgroup priority selection circuit 26, except that a change is effected after each device controller 14 is selected.

While the above provides a full and complete disclosure of the invention, it should be obvious to those skilled in this art that modifications and variations of the disclosed invention can readily be implemented without departing from the scope of the invention as defined by the appended claims.

CLAIMS

1. In a data processing system of the type including a processor unit (12), a plurality of peripheral units (14), and a bus means (16) inter-
5 connecting the processing unit and the peripheral units for data communication therebetween, apparatus for selecting one of two or more peripheral units desiring data communication with the processor unit, the apparatus being characterised by:

10 means for polling the peripheral units to determine which of the peripheral units desire said data communication;

means responsive to the polling means for selecting a one of the peripheral units desiring
15 data communication for data communication according to a predetermined priority schedule; and

means coupled to the selecting means for changing the predetermined priority schedule.

2. Modifiable priority selection apparatus
20 for determining which of a number of peripheral devices, requesting to do so, shall communicate with the processor unit of a data processing system, the modifiable priority selection apparatus being characterised by:

25 register means settable to a state representative of one of a number of predetermined priority selection schedules;



logic means coupled to the register means and responsive to the number of peripheral devices to
30 provide an indication of a selective one of the peripheral units according to the one of the number of predetermined priority selection schedules represented by the register means; and

means for modifying the content of the register
35 means.

3. A modifiable priority selection apparatus as claimed in claim 2, wherein the modifying means includes means for modifying the register means according to a predetermined pattern.

- 5 4. A modifiable priority selection apparatus as claimed in claim 3, there being a number of predetermined patterns, and means presettable for selecting which of the predetermined patterns will be used.

.....



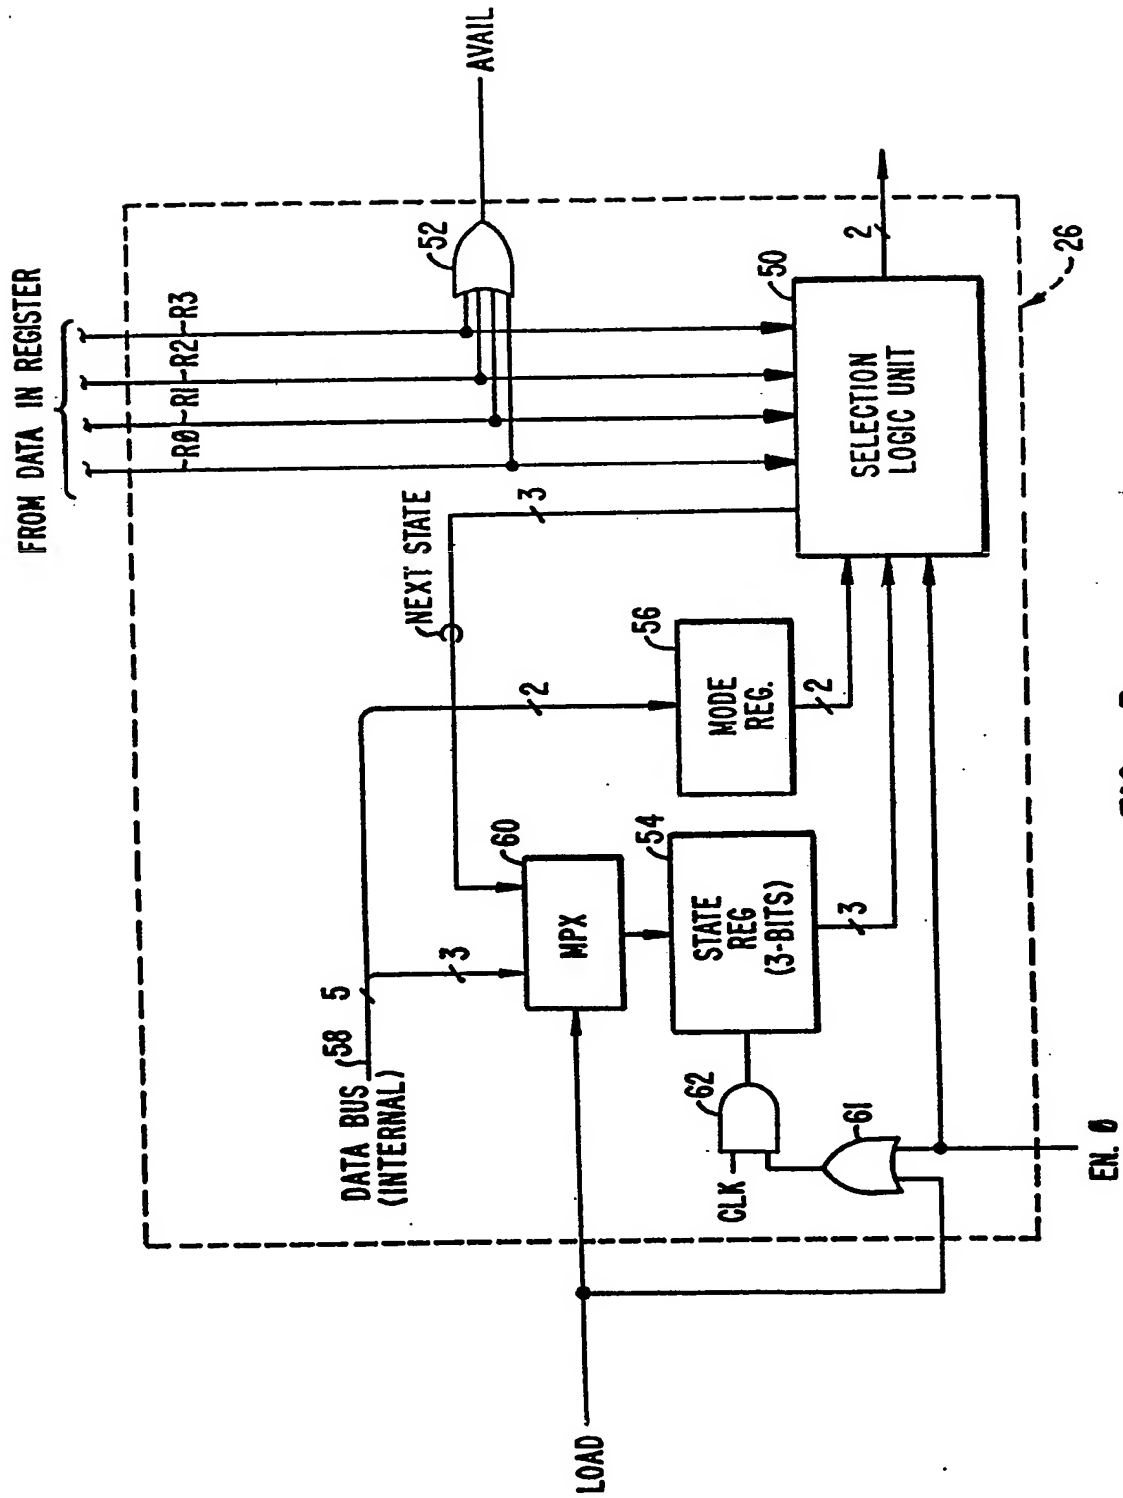


FIG. 3.